

CLAIMS

1. A method for handling errors and events arising in a storage shelf containing storage devices interconnected through path controller cards to storage-shelf routers contained on storage-shelf-router cards, the method comprising:

5 detecting an error or event;

diagnosing the error or event;

for an error or event remedied by automatically initiated replacement of a storage-shelf-router card, failing over path controller cards primarily managed by one or more storage-shelf routers on the storage-shelf-router card to be replaced to one or more different storage-shelf routers on surviving storage-shelf-router cards; and

10 for an error or event within a mass-storage device or path-controller card, and for other errors and events that are configured for external management, reporting and logging the error or event for handling by an entity external to the storage shelf.

15 2. The method of claim 1 wherein the detected error is an external FC link failure ("EFCLF"), detected by events including link-down events and occurrence of cyclic-redundancy-check errors, diagnosed by loop-back tests following isolation of an internal FC mini-hub on a storage-router card, and remedied by automatic or directed fail over of a storage-router card.

20

3. The method of claim 1 wherein the detected error is an internal link failure ("ILF"), detected by events including link-down events and occurrence of cyclic-redundancy-check errors, diagnosed by loop-back tests following isolation of an internal FC mini-hub on a storage-router card, and remedied by automatic or directed fail over of a storage-router card.

25

4. The method of claim 1 wherein the detected error is an inter-card port failure ("ICPF"), detected by loss of heartbeat, diagnosed by an alternative path-controller-card-mailbox-based communication between storage-router cards and pad tests, and remedied by fail over of a storage-router card.

30

5. The method of claim 1 wherein the detected error is an inter-card link failure ("ICLF") (""), detected by loss of heartbeat, diagnosed by an alternative path-controller-card-mailbox-based communication between storage-router cards and pad tests, and remedied by fail over of a storage-router card.

5

6. The method of claim 1 wherein the detected error is an SATA port failure ("SPF"), detected by events including link-down events and occurrence of cyclic-redundancy-check errors, diagnosed by pad tests and loop back tests, and remedied by fail-over a storage-router card of event notification.

10

7. The method of claim 1 wherein the detected error is a management port failure ("MPF"), detected by periodic access of path-controller-microprocessor registers, diagnosed by internal storage router loopback, and remedied by reporting of the error to an external processing entity.

15

8. The method of claim 1 wherein the detected error is an uncontrolled critical failure, detected by loss of heartbeat, and remedied by fail-over of a storage-router card.

9. The method of claim 1 wherein the detected error is a controlled critical failure ("CCF"), detected by loss of heartbeat, and remedied by fail-over of a storage-router card.

20

10. The method of claim 1 wherein the detected error is a peer field replaceable unit removal ("PFR"), detected by de-assertion of a PEER_PRESENT signal, diagnosed by link testing, and remedied by fail-over of a storage-router card.

25

11. The method of claim 1 wherein the detected error is a field-replaceable-unit insertion fail back ("FBE"), detected by PEER-PRESENT signal assertion, diagnosed by a rendezvous protocol between storage-router cards, and remedied by a fail-back of path-controller cards to a newly inserted storage-router card.

30

12. The method of claim 1 wherein the detected error is a SATA management link failure ("MLF"), detected by periodic access of path-controller-microprocessor registers, diagnosed by internal storage router loopback, and remedied by reporting of the error to an external processing entity.

5

13. The method of claim 1 wherein the detected error is a SATA drive failure ("SDF"), detected by a failure of a command directed to a disk, and remedied by reporting of the SDF event to an external processing entity.

10 15. The method of claim 1 wherein the detected error is a drive-field-replaceable-unit removal ("FRE"), detected by de-assertion of a FRU_PRESENT signal and remedied by loop initialization and reporting of the FRE to an external processing entity.

15 16. The method of claim 1 wherein the detected error is a drive- field-replaceable-unit insertion ("FIE"), detected by, diagnosed by, and remedied assertion of a FRU_PRESENT signal and remedied by disk initialization and loop initialization and reporting of the FIE to an external processing entity.

20 17. The method of claim 1 further including, for an error or event remedied by failing over a single path-controller card, failing over the path-controller card from a first storage router on a first storage-router card to a second storage router on a second storage router card.

18. A method for replacing a storage-router card in a storage shelf, the method comprising:

25 failure of a first storage-router card;
detection of the failure of the first storage-router card by a second storage-router card;
carrying out a local path failover by the second storage-router card;
replacing the first storage-router card;
detection of the replacement of the first storage-router card by the second storage-
30 router card;
boot-up and initialization of the replaced first storage-router card;

when the replaced first storage-router card properly initializes, carrying out a local path fail back to the replaced first storage-router card; and

when the replaced first storage-router card does not properly boot, carrying out a local path fail over to the replaced first storage-router card and replacing the second storage-router card.

19. The method for replacing a storage-router card in a storage shelf of claim 18 further comprising:

failure of a first storage-router card;

10 detection of the failure of the first storage-router card by a second storage-router card; carrying out a local path failover by the second storage-router card;

asserting a hard reset by the second storage-router card on the first storage-router card following carrying out a local path failover by the second storage-router card and preceding replacing the first storage-router card; and

15 de-assertion of the hard reset by the second storage-router card following detection of the replacement of the first storage-router card by the second storage-router card and preceding boot-up and initialization of the replaced first storage-router card.

20. A storage shelf within a storage system comprising:

20 a storage shelf having

at least two storage-routers included in at least two storage-router cards,

a high-bandwidth interconnection between the two storage-router cards,

a number of storage devices, each storage device interconnected to a path-controller card, the path controller card interconnected to two storage routers; and

25 failure domains, a failure domain comprising one of a storage-router card, a path-controller card and associated storage device, and components exterior to the storage-router-card failure domain, the path-controller-card-and-associated-storage-device failure domain, and interconnections between storage-router cards and storage-router-cards and path-controllers.